IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:	Confirmation No. 5905
Mark A. Kampe et al.	Art Unit: 2194
Serial No.: 09/846,254	Examiner: Andy Ho
Filed: May 2, 2001	Customer No. **32658**
For: **CLUSTER EVENT SERVICE METHOD AND SYSTEM**	Docket No. P5088

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF UNDER 37 CFR § 41.37

I. Real Party in Interest

Sun Microsystems, Inc.
4120 Network Circle
Santa Clara, CA 95054
USA

II. Related Appeals and Interferences

No other appeals or interferences are currently known to Appellants that will directly affect, be directly affected by, or have a bearing on the decision to be rendered by the Board of Patent Appeals and Interferences in the present appeal.

III. Status of Claims

Claims 1-4, 6-12, 14-22, 33-40, 43, 62-65, 67-80, 85, 86, 91, and 94 are pending in the application, with claims 5, 13, 23-32, 41, 42, 44-61, 66, 81-84, 87-90, 92, and 93 being cancelled. No claims have been allowed, and all pending claims stand rejected under 35 U.S.C. §103. The rejection of claims 1-4, 6-12, 14-22, 33-40, 43, 62-65, 67-80, 85, 86, 91, and 94 is the subject of this appeal.

IV. Status of Amendments

No claim amendments were filed subsequent to a final rejection.

Claims 1-4, 6-12, 14-22, 33-40, 43, 62-65, 67-80, 85, 86, 91, and 94 are provided in the attached Claims Appendix.

V. Summary of Claimed Subject Matter

Claims 1, 12, 33, 62, 70, 74, 80, and 91 are independent claims that are being appealed.

Claim 1 is directed to a network having nodes connected by a digital communication link. Figure 1 of Appellants' specification shows an exemplary network 100 of the invention with nodes 101(a)-101(d) connected by communication link 107 and the description of this network 100 begins at paragraph [0011]. The network of claim 1 includes "an event channel adapted to transfer an event between a publisher node and a subscriber node within said network over the communication link." Elements 106(a) and 106(b) of Figure 1 are representative of such an event channel with each of the nodes 101(a)-101(d) being a publisher node and/or a subscriber node. An "event channel" is described in paragraph [0016] as "a communication pipe that allows multiple suppliers to communicate with multiple consumers asynchronously." Events and event channels are discussed in paragraph [0015], and a representative "event" is shown with element 300 of Figure 3, and the event is described in paragraph [0028] as being used by applications to exchange information, with the event 300 including a data field 302 and pattern fields 304, 306, 308 that are used to identify differing types of events 300 and to allow a subscribing application to properly filter events (e.g., to receive only a subset of events 300 published on a particular event channel).

The network of claim 1 further includes "a filter on said subscriber node" that processes events published on the event channel to identify matching events "wherein said matching event includes at least one pattern field that matches a filter field within said filter." Further, claim 1 calls for "an application on said subscriber node to receive the matching event and the "application defines said filter fields within said filter and opens said event channel at said subscriber node." Figure 1 shows applications 102(a)-102(d) at each node, and Figure 2

illustrates a node 201 in more detail that may correspond to nodes 101(a)-101(d) of Figure 1. As shown, each application 202(1), 202(2), and 202(3) includes a filter 203(1), 203(2), and 203(3) that communicates with an event server. These components are described in detail in paragraphs [0019] and [0022], with a subscriber application acting to register a filter "to identify the subset of events that it wishes to receive" that is typically a subset of all event published on an event channel. The filtering mechanism of these application-defined filters "may be based on pattern matching on the events published on event channel." Paragraph [0029] goes on to state "the filter is composed of three patterns that will be matched against the three patterns of every event published on the event channel. Those events that match the filter patterns may be delivered to the client application that subscribed to the event." Additionally, Figure 4 illustrates a flowchart depicting how information is received in a network such as that claimed in claim 1 and describes how event channels are opened by applications and filters are defined by applications for use on a subscriber node.

Independent claim 12 is directed to a node within a network. This node may be nodes 102(a)-102(d) of Figure 1 or node 201 of Figure 2, for example. The node includes an application running on the node, such as applications 102(a)-102(d) shown in Figure 1 or applications 202(1)-202(3) of node 201 in Figure 2. The node includes a filter which was described in detail with reference to claim 1, and this description is applicable to the filter of claim 2 (e.g., the filter is "assigned by said application"). The node further includes "an event server" with "an event control block to subscribe to said event channel for said application." Representative event servers are shown with elements 105(a)-105(d) in the nodes of Figure 1 and with element 205 of the node of Figure 2. Further, a detailed event server 700 is shown in Figure 7 and includes event control blocks (ECB) 701(a)-701(c). Operation of event servers of the invention are discussed in paragraphs [0014], [0016], [0020], [0021], and [0039]-[0046]. Claim 12 also calls for the event to be placed in a queue on the node by the event server prior to use by the application, which is described at paragraph [0042] as being performed by event control blocks 701(a)-701(c).

Independent claim 33 is directed to a method for receiving information at a node. The steps of claim 33 have been described in the discussion of the operation of the components of the

network of claim 1 and the node of claim 12, and those summaries are applicable to claim 33. Further, the flowchart provided in Figure 4 is applicable to claim 4 with block 406 showing opening an event channel at the node, with block 410 showing using an application on the node to assign a filter to the event channel, and with block 412 showing the receiving of an event according to the filter. The processes of receiving information at a node are described in more detail with reference to Figure 4 in paragraphs [0031]-[0033], and as part of the description of node 201 of Figure 2 in paragraphs [0019]-[0023] and [0029]. Independent claim 80 is similar to that of claim 33 in that it is directed to receiving information at a node and includes similar limitations. Hence, the summary of the invention of claim 33 is applicable to the invention of claim 80.

Claim 62 is directed to a method for declaring a node to an event server. The method includes "providing an event server on a node of computer network" such as event servers 105(a)-105(d) or 205 on nodes 101(a)-101(d) or 201 in Figures 1 and 2. Also, Figure 8 illustrates a flowchart applicable to claim 62 and shows the process for declaring a new subscriber and/or a new publisher node. The method continues with "granting the event server access to an event channel provided on a digital data communications link" with the event channel already being described with reference to claims 1 and 12 and paragraph [0072] discussing granting access with reference to block 804 of Figure 8. The method continues with "creating a naming context for said event channel" as shown in block 802 of Figure 8, which ids discussed at the start of paragraph [0072], with each channel being able register unique names with a naming service. The method includes "updating an event control block" in the event server and the access corresponds "to an application running on the node." This step may include one or more blocks 806-814 as described in paragraphs [0072]-[0074] with reference to the ECBs 701 of Figure 7. The method further includes using the event server to send "a filter control message over the event channel to another event server at another node" which is shown at block 824 of Figure 8. Paragraphs [0076] and [0077] describe the use of filtering control messages to control message transmissions from event servers on particular event channels.

Such filtering may be thought of as "source filtering" and independent claim 74 is similar to claim 62 in that it is directed to a method of "source filtering at an event server on a publisher

node within a network." The method includes "sending a filter control message to said publisher node;" "marking a remote event control block object in an event control block according to said filter control message;" and "filtering events from said event control block." As with claim 62, the steps of the method 74 are found in Figure 8 (such as block 824) and are also seen in Figure 13 with reference to block 1316.

The method of claim 70 is for implementing distributed filtering at a node (such as nodes 101(a)-101(d) or 201 of Figures 1 and 2). The distributed filtering techniques of the invention are described in paragraphs [0086]-[0105] with reference to Figures 9-13. The method calls for "building a filter from a plurality of search trees" with such trees shown Figures 9 and 10 (and described in corresponding text) and, as discussed above, the application typically functions to define one or more filters for determining what type of information it is to receive over an event channel and building described at least in paragraph [0093]. The method continues with "receiving an event at said node" with "an event" being shown in Figure 3 and its receipt discussed with reference to claims 1 and 12. The method further includes "selecting a search tree from said filter." The flowchart of Figure 13 is useful for discussing claim 70 and is described beginning at paragraph [0104] with the event protocol module 704 of event server 700 of Figure 7 being used to select an appropriate tree at block 1302 of Figure 13, such as "the one that has the smallest height in the lexicographic tree and does not start with a null node value." The method of claim 70 further includes comparing the event with the search tree, such as with one or more components of an event server (such as that shown in Figure 7). Independent claim 91 is directed to a computer program product with limitations similar to those of claim 70, and the summary of the invention of claim 70 is applicable to claim 91.

VI. Grounds of Rejection to be Reviewed on Appeal

Claims 1-4, 6-12, 14-22, 33-40, 43, 62-65, 67-69, 74, 75, 78-80, 85, 86, and 94 stand rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Pat. No. 6,477,585 ("Cohen") in view of U.S. 6,658,487 ("Smith").

Claims 70-73, 76, 77, and 91 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Cohen in view of Smith and further in view of U.S. Pat. No. 6,314,533 ("Novik").

VII. Argument

Rejection of Claims 1-4, 6-12, 14-22, 33-40, 43, 62-65, 67-69, 74, 75, 78-80, 85, 86, and 94 Under 35 U.S.C. §103(a) Based on Cohen and Smith is Improper

In the final Office Action of March 8, 2006, claims 1-4, 6-12, 14-22, 33-40, 43, 62-65, 67-69, 74, 75, 78-80, 85, 86, and 94 were rejected under 35 U.S.C. §103(a) as being unpatentable based on Cohen in view of Smith. This rejection is traversed based on the following remarks, and Appellants request that the rejection be reversed as not properly supported.

Initially, as background information for the review board, on August 4, 2005, the Appellants filed a Notice of Appeal including a Pre-Appeal Brief Request for Review. In response, prosecution was reopened with the Examiner performing an additional search resulting in the rejection of all claims on new grounds based on previously-cited U.S. Pat. No. 6,477,585 ("Cohen") and on newly-cited U.S. Pat. No. 6,658,487 ("Smith"). Cohen had previously been used as an anticipatory reference but now is supplemented with the teaching of Smith and the rejections are now all obviousness rejections.

Claim 1 calls for a filter to be provided "on said subscriber nodes," and this filter acts to "**process a plurality of events published on said event channel to identify said event as a matching event**." With this configuration, the subscriber node does the filtering of events published or made available on a linked event channel, i.e., the network uses subscriber-side filtering. In contrast, Cohen teaches supplier or publisher-side filtering. Hence, the network of claim 1 is not shown or suggested by Cohen.

As discussed in the several of Appellants' Amendments, with regard to claim 1, the Office Action cites Cohen at col. 5, lines 48-49 for teaching the event channel of claim 1 and at col. 6, line 7 (consumer-side EMS filter) and col. 6, lines 19-22 for showing a filter to identify an event on the subscriber node. Appellants disagree with this construction of Cohen. At col. 5, lines 55-61 with reference to Figures 2 and 3, Cohen makes it clear that its event distribution method involves providing a single host computer running an event management system (EMS

22), i.e., the supplier or publisher that performs the filtering. According to Cohen, clients must subscribe to the EMS 22 and also define filters that are stored in a filter database 46 at the device hosting the EMS 22 (i.e., not on the event consumers 26a-26n). Also, with reference to Figure 3, the event channel is shown to be part of the EMS 22. Based on these arguments, Cohen fails to shown "a filter on said subscriber node" because as can be seen in Figure 3 the event consumers 26 are remote to the EMS 22 which stores the filters in database 46.

The September 21, 2005 and more recent March 8, 2006 Office Actions indicate that the Examiner agrees with this argument that Cohen fails to teach the filter feature of claim 1. Smith is then cited at lines 36-40, col. 3, lines 58-60, col. 3, and Figs. 3 and 4 for teaching "an event filter is located on the same node as the event subscriber." Turning to Smith's teaching, it can be seen in Figures 3, 4, 6, and 7 that the described distributed computing system includes an element labeled a "filter" that is associated with objects (such as client and server objects 10, 12 of Fig. 3). However, the Smith "filter" does not teach the filter of claim 1 as the claimed filter is provided "to process a plurality of events published on said event channel to identify said event as a matching event" and such matching event "includes at least one pattern field that matches a filter field within said filter." Hence, the filter on the subscriber node of claim 1 identifies matching events based on a pattern in the event and a filter field in the filter.

The Smith "filter" does not perform any such matching function and the "events" of Smith are not described as having a pattern field. As a result, Smith does not overcome the deficiencies of Cohen because it does not teach an active filtering or matching mechanism at each subscriber node. This can be seen by studying Smith from col. 4, line 6 to line 46. Smith teaches that its "filters" do not perform any matching to identify events that an application receives (see the last element of claim 1 where matching events are received by the application). Instead, Smith teaches that the filters intercept outgoing messages and incoming messages and pass these intercepted messages to an "event collection mechanism." There is no matching or filtering but instead all events are apparently passed on to the event collection mechanism (element 14 in the figures). With reference to Figures 6 and 7, the Smith filters are described as attaching keys to messages prior to passing on to a receiving filter, but there is no discussion of a receiving filter matching or filtering received messages with a pattern it stores to identify matching messages for receipt by an associated object.

However, it should be understood that Smith does teach that filtering is performed in its distributed system – just not at the filters associated with the distributed objects. For example, at col. 6, lines 14-45, Smith states that "necessary filtering is carried out by the respective Event Dispatcher 54 which passes on only that information which has been requested by its individual Visualizer Application." Hence, Smith teaches that actual filtering is performed by an event dispatch mechanism 16 as shown in Figure 9 rather than by the filters associated with the objects. For these reasons, Smith fails to overcome the deficiencies of Cohen. Further, the combination of the two references' teachings would not result in the claimed filter because both teach filtering of messages at a central location and not at the node running an application.

The Response to Arguments of the March 8, 2006 Office Action states that Smith is only cited for teaching a filter located on an application node and, apparently, not the particulars of such a filter. However, as discussed above, the Smith filtering does not occur at the application or subscriber node. Hence, if one were to modify Cohen with the teaching of Smith, the invention of claim 1 would not be achieved as both references teach centralized filtering. The only motivation to perform the actual filtering of events or messages is found in the Appellants' specification, and it is impermissible to use Appellants' own teaching to provide the sole motivation to modify a primary reference. The teaching with Smith may possibly teach the usefulness of intercepting messages at a subscriber and transmitting them to a central device for filtering rather than the publishing device as in Cohen but would not teach the network of claim 1 with a filter on the subscriber node that is used to identify matching events which the application receives (without passing the messages to another device for further processing/filtering).

Further, Cohen fails to teach "an application...opens said event channel at said subscriber node." The Response to Arguments in all prior Office Actions fail to address this argument for allowing claim 1 over Cohen provided in the last response (and the September 21, 2005 Office Action simply restates its assertions that Cohen teaches this limitation). The Office Action cites Cohen at lines 48-49 of column 5 for providing this teaching. Cohen, at this citation, states "Communications through the event channel are "asynchronous" in that they may be provided to the event consumers at any time." Cohen does NOT teach that an application at the subscriber node that defines the filter and its fields also acts to open an event channel provided between the

publisher and the subscriber nodes. If the event consumers of Cohen are taken to be the subscriber nodes, there is no discussion in Cohen that an application on these nodes acts to open an event channel. From col. 5, lines 14-37, it appears that communications between the EMS/event suppliers and the event consumers is controlled by the EMS. For this additional reason, Cohen fails to teach or suggest each and every limitation of claim 1, and as noted by the Examiner, Smith is only "used to teach the advantages of having a filter locating [sic] at the subscriber" and not for overcoming this deficiency in Cohen.

Claims 2-4 and 6-11 depend from claim 1 and are believed allowable as depending from an allowable base claim. Claim 94 also depends from claim 1 and is believed allowable as Cohen fails to teach a plurality of subscriber nodes each including **a filter defined by an application on the node**, opening an event channel over a communication link to each such node, and using the filter at each node to identify matching events for receipt by the application. Additionally, it should be noted that Smith does not teach that its "filters" are defined by the associated objects, and Appellants requested in their last Amendment that claim 94 be found allowable or additional references be cited showing this additional limitation of claim 94. The Examiner has not responded to this request or addressed the specific arguments regarding the additional limitations of claim 94 that Appellants assert are not shown in Cohen or in Smith.

Regarding independent claim 12, the Office Action relies on Cohen and Smith to reject the claim in a manner similar to that of claim 1. Therefore, the reasons for allowing claim 1 over Cohen and Smith are applicable to claim 12. Additionally, Cohen fails to teach a queue on the same node that assigns the filter and receives and uses matching events. In contrast, the queue 47 is shown to be part of the EMS 22 and is placed on single host within a network as shown in Figures 2 and 3 (e.g., not on the consumer nodes 26). This additional reason for allowing claim 12 was provided in the last response and the Pre-appeal Brief Request for Review, but the Examiner did not address the argument in the Response to Arguments in the prior two Office Actions (i.e., the event log is not on the node of the application). For this additional reason, the rejection of claim 12 based on Cohen is not proper and should be withdrawn. Smith does not overcome these additional deficiencies of Cohen with respect to claim 12. Specifically, Smith does not teach that the matching events are placed on a queue on said node by its "filter" elements and the Examiner admittedly did not cite Smith for overcoming this problem with

Cohen. As a result, the combination of Cohen and Smith does not support a rejection of claim 12, and claims 12 and claims 14-22, which depend from claim 12, are believed allowable.

Independent claim 33 calls for opening an event channel at a node that provides a shared communication path on a communication link and to subscribing to receive events at the node over the event channel. Cohen fails to teach these features as it describes (as discussed with reference to claim 1) running an EMS on a single node and then distributing events to specific nodes after filtering on the EMS node. The method of claim 33 is very different in that it supports fully asynchronous communication over the event channel without requiring an event publisher to provide addresses of receiving nodes as opposed to the API 32 and service 22 of Cohen as described with reference to lines 43-46, col. 5.

The method of claim 33 includes running an application on the node, receiving and processing an event at the node over the event channel, and then when a match is determined "at said node" passing the received event to the application on the node. Distribution out of the node is not required after filtering as is the case in the Cohen method. For these reasons, claims 33 and claims 34-40 and 43, which depend from claim 33, are believed allowable over Cohen.

Smith fails to overcome these deficiencies of Cohen for the reasons provided for claim 1. Specifically, Smith's "filters" do not perform receiving an event over an event channel as called for in claim 33. Further, as discussed with reference to claim 1, Smith fails to teach finding "a match according to said filter" and when a match is found "passing the received event to the application on the node." In other words, where does Smith teach that it passes matches to an object? In contrast, it shows sending all events received by a "filter" on to an associated object and not really "filtering" anything at the subscriber node. For these reasons, the combined teaching of Cohen and Smith fails to support a rejection of claim 33. The Examiner has not addressed these separate arguments for allowance of claim 33 in the Response to Arguments of the March 8, 2006 Office Action and has rejected claim 33 "for the same reasons as claim 1" with no explanation of how Cohen teaches the above-discussed differing features of claim 33 relative to claim 1. Claims 33 and claims 34-40 and 43, which depend from claim 33, are believed in condition for allowance.

Independent claim 62 was rejected in the final Office Action for the same reasons as provided for rejecting claim 1 (and its dependent claim 15), and the reasons provided for

allowing claim 1 over Cohen and Smith are applicable to claim 62. Further, Cohen and Smith fail to teach or suggest granting access to an event channel on a communication link and associating such access or permission to an application running on a node network. The Office Action cites Cohen at line 59, col. 12, line 12, col. 14, and lines 34-35, col. 5 for showing these elements not presented in claim 1. Appellants could find no teaching of this limitation and particularly, of associating such access to an application running on a node network in Cohen. Further, Cohen fails to show creating a name context for the event channel as called for in claim 62, with the cited portion of Cohen at col. 11, line 28 simply referring to "the EMS event channel" but providing no teaching of providing a name context for a created event channel that, as will be appreciated, can later be used for providing events to subscribers of a particularly named event channel. Hence, Cohen and Smith do not support a rejection of claim 62 or claims 63-69, which depend from claim 62, and these claims are believed in condition for allowance. The final Office Action provided no response to these arguments regarding claim 62.

Regarding independent claim 74, the Office Action again states that claim 74 is the same method as claims 1, 13, 14, and 62 and rejects it for the same reasons as these claims. However, claim 74 includes differing limitations not included in claims 1, 14, and 62 (with claim 13 being cancelled. Specifically, claim 74 calls for "marking a remote event control block object in an event control block according to said filter control message," and none of the claims mentioned by the Examiner include this limitation. For example, claim 1 does not discuss a filter control message, an event control block, or a remote event control block object (or marking such an object). Claim 14 states "wherein said event server further includes an event control manager to control said event control block" and this language does not include the limitations of claim 74. Claim 62 discusses an event control block and sending a filter control message but does not including "marking a remote event control block object in an event control block according to said filter control message." A proper obviousness rejection of claim 74 requires a separate rejection indicating where each of its elements are shown or suggested in Cohen and/or Smith. This has not been provided in any of the Office Actions to date. Hence, the Examiner has failed to make out a proper case of obviousness because the Examiner has not provided explicit citations to Cohen or Smith where each and every limitation in the claim is shown or made

obvious. As a result, claim 74 and claims 75, 78, and 79, which depend from claim 74, are believed in condition for allowance.

Independent claim 80 was rejected in the Office Action for the reasons provided for rejecting claim 1, and hence, the reasons provided for allowing claim 1 over Cohen and Smith are believed applicable to claim 80. Specifically, Cohen fails to teach using a client application for opening an event channel on the same node as is running the application and receiving and filtering events on the channel with a filter on the application's node. Further, Cohen fails to teach opening such an event channel in read and write modes as called for in claim 80. The cited col. 9, lines 41-62 do not mention opening an event channel in a read mode or in a write mode or that such opening can be done by a client application on a node of a network. Based on these arguments, claim 80 and claims 85 and 86, which depend from claim 80 are not shown or suggested by Cohen, and the rejection of these claims should be withdrawn. The final Office Action does not address these arguments for allowing claim 80.

Rejection of Claims 70-73, 76, 77, and 91 Under 35 U.S.C. §103(a) Based on Cohen, Smith, and Novik is Improper

Additionally, in the Office Action of March 8, 2006, claims 70-73, 76, 77, and 91 were rejected under 35 U.S.C. §103(a) as being unpatentable based on Cohen in view of Smith and further in view of Novik. This rejection is traversed based on the following remarks, and Appellants request that the rejection be reversed.

Referring to independent claim 70, the Office Action states that Cohen fails to teach building its filters from a "binary tree" but cites Novik at col. 2, lines 56-59, Figure 6, and at col. 14, lines 40-53 for providing teaching building filters from "search trees" (as called for in claim 70). However, at this citation, Novik states "Preferably, the filtering of events would be performed at the event provider itself, such that any events that are not requested by a subscriber would be discarded at the event provider." There is no teaching at this citation of building a filter from a plurality of search trees, of selecting a search tree from said filter, and comparing said event with said search tree as called for in claim 1.

Further, Novik teaches similarly to Cohen that filtering is performed at the event supplier or publisher. In contrast, claim 70 calls for the building, selecting, and use of the filter to be

performed at the node that is also used for "receiving an event at said node." Hence, the filtering (and its construction) are performed at the event consumer or subscriber rather than at the event supplier or provider node as taught by both Cohen and Novik. Smith as discussed with reference to claim 1 also fails to teach "filtering" at a node associated with an object with the Smith "filters" simply passing all events through with copies being passed to an event collection device. Since these references fail to teach or suggest each and every limitation of claim 70 and actually teach away from its limitations, claim 70 is not made obvious by the combined teachings of these two references. The final Office Action fails to address or rebut these arguments made by Appellants for the allowance of claim 70 and the failings of Novik to overcome the admitted deficiencies of Cohen.

Claims 71-73 depend from claim 70 and are believed allowable for at least the reasons provided for allowing claim 70.

Claims 76 and 77 depend from claim 74 and are believed allowable as depending from an allowable base claim. Further, Novik fails to overcome the deficiencies noted with reference to claim 74 in Cohen and Smith.

Independent claim 91 is directed to a computer program product with limitations similar to that of claim 70. The reasons provided above for allowing claim 70 over Cohen and Novik are believed applicable to claim 91.

Conclusion

In view of all of the above, all the pending claims are believed to be allowable and the case in condition for allowance. Appellants respectfully request that the Examiner's rejections based on 35 U.S.C. §103 be reversed for all the pending claims.

Respectfully submitted,

Date: July 25, 2006

Kent A. Lembke, Reg. No. 44,866
HOGAN & HARTSON LLP
One Tabor Center
1200 17th Street, Suite 1500
Denver, Colorado 80202
Phone: (720) 406-5378
Fax: (720) 406-5301

VIII. CLAIMS APPENDIX

1. A network having a plurality of nodes connected by a digital data communication link, comprising:

an event channel adapted to transfer an event between a publisher node and a subscriber node within said network over the communication link;

a filter on said subscriber node to process a plurality of events published on said event channel to identify said event as a matching event, wherein said matching event includes at least one pattern field that matches a filter field within said filter; and

an application on said subscriber node to receive said matching event, wherein said application defines said filter fields within said filter and opens said event channel at said subscriber node.

2. The network of claim 1, further comprising an event server on said subscriber node, said event server adapted to receive said event and to pass said received event to said filter from said event channel.

3. The network of claim 2, wherein said event server exchanges information with another event server on another one of the nodes of the network.

4. The network of claim 2, wherein said application opens said event channel through said event server.

6. The network of claim 1, wherein said event further includes a data field.

7. The network of claim 1, wherein said event channel has a unique name.

8. The network of claim 7, wherein said unique name is registered in a naming service within said network.

9. The network of claim 2, wherein said publisher node has a configuration, said configuration being known to said event server on said subscriber node.

10. The network of claim 1, further comprising an event server on said publisher node, wherein said event server publishes said event on said event channel.

11. The network of claim 10, wherein said subscriber node has a configuration, said configuration being known to said event server on said publisher node.

12. A node within a network to exchange information, comprising:

an application running on the node;

an event server adapted to receive events from an event channel on a communication link, wherein said event server includes an event control block to subscribe to said event channel for said application; and

a filter to identify matching ones of said events for use by said application, wherein said filter is assigned by said application, said event includes at least one pattern field, said at least one pattern field matches at least one filter field within said filter, and said event is placed in a queue on said node by said event server prior to the use by said application.

14. The node of claim 12, wherein said event server further includes an event control manager to control said event control block.

15. The node of claim 14, wherein said event control manager updates said event control block.

16. The node of claim 14, wherein said event control manager detects an overload condition within said event control block.

17. The node of claim 14, wherein said event control manager controls a configuration of said event control block.

18. The node of claim 12, wherein said event server further includes an event protocol module to manage network connections to said event control block.

19. The node of claim 12, wherein said event control block includes a remote event control block that correlates to an event control block.

20. The node of claim 12, wherein said event server includes an event channel descriptor to access said event control block for said client application.

21. The node of claim 12, further comprising an event application program interface to publish and subscribe to said event channel.

22. The node of claim 12, wherein said event is processed by said application.

33. A method for receiving information at a node, comprising:

opening an event channel at said node, said event channel providing a shared communication path on a digital data communication link with other of said nodes;

subscribing to receive events at the node over via the event channel;

with an application running on the node, assigning a filter to said event channel;

receiving an event on said event channel;

processing said event at said node to determine whether the received event is a match according to said filter; and

when determined a match, passing the received event to the application on the node.

34. The method of claim 33, further comprising registering a function indicating said opened event channel.

35. The method of claim 33, further comprising publishing said event on said event channel.

36. The method of claim 33, further comprising dispatching a callback responding to said event.

37. The method of claim 33, further comprising creating said event channel by operation of an event server running on the node.

38. The method of claim 33, further comprising filtering said event by said filter.

39. The method of claim 38, wherein said filtering includes matching a pattern within said event with a filter pattern within said filter.

40. The method of claim 33, further comprising storing said event at an event control block.

43. The method of claim 33, wherein said subscribing includes invoking an event control block.

62. A method for declaring a node to an event server, comprising:

providing an event server on a node of a computer network;

granting the event server access to an event channel provided on a digital data communications link;

creating a naming context for said event channel;

updating an event control block in said event server reflecting said granted access, wherein the granted access corresponds to an application running on the node; and

with the event server, sending a filter control message over the event channel to another event server at another node.

63. The method of claim 62, further comprising allocating said event control block.

64. The method of claim 62, further comprising finding said event control block on said event server.

65. The method of claim 62, further comprising getting a naming context for said event channel.

67. The method of claim 62, further comprising unlocking said event control block.

68. The method of claim 62, further comprising changing an access permission to said event channel.

69. The method of claim 62, further comprising returning to an application at said node.

70. A method for implementing distributed filtering at a node, comprising:

building a filter from a plurality of search trees;

receiving an event at said node;

selecting a search tree from said filter; and

comparing said event with said search tree.

71. The method of claim 70, further comprising building said plurality of search trees.

72. The method of claim 71, further comprising placing heads from said plurality of search trees within said filter.

73. The method of claim 70, further comprising modifying said search trees.

74. A method for source filtering at an event server on a publisher node within a network, comprising:

sending a filter control message to said publisher node;

marking a remote event control block object in an event control block according to said filter control message; and

filtering events from said event control block.

75. The method of claim 74, further comprising building said filter control message.

76. The method of claim 75, further comprising selecting a filter search tree.

77. The method of claim 76, further comprising modifying said filter search tree.

78. The method of claim 74, further comprising changing access permissions to a remote event control block and re-sending said filter control message.

79. The method of claim 74, further comprising unmarking said remote event control block object.

80. A method for receiving information at a node, comprising:

opening an event channel with a client application running on said node, the client application opening said event channel in a write mode or a read mode, wherein the client application can publish events over the event channel in said write mode and can receive events published on the event channel in said read mode;

receiving an event from said event channel at said node;

assigning a filter to said event channel by said client application running on said node; and

filtering said event from said event channel with said filter at said node.

85. The method of claim 80, further comprising queuing said event in an event control block at said node corresponding to said application.

86. The method of claim 80, further comprising dropping said event during an overload condition at said node.

91. A computer program product comprising a computer useable medium having computer readable code embodied therein for implementing distributed filtering at a node, the computer program product adapted when run on a computer to effect steps, including:

building a filter from a plurality of search trees;

receiving an event at said node;

selecting a search tree from said filter; and

comparing said event with said search tree.

94. The network of claim 1, further comprising a plurality of additional publishers nodes linked to the event channel transferring publishing events on the event channel and a plurality of additional subscriber nodes linked to the event channel, each of the additional subscriber nodes comprising a filter defined by an application on the additional subscriber node with filter fields for use in processing said published events to identify matching events based on the filter fields.

IX. EVIDENCE APPENDIX

No copies of evidence are required with this Appeal Brief. Appellants have not relied upon any evidence submitted under 37 C.F.R. §§ 1.130, 1.131, or 1.132.

X. RELATED PROCEEDINGS APPENDIX

There are no copies of decisions rendered by a court or the Board to provide with this Appeal as there are no related proceedings.